

WHAT IS CLAIMED IS:

1. A block-level shared network storage system, comprising:
a storage server comprising an array of disk drives, and comprising a processor that runs a device driver to provide block-level access to data stored
5 on the array of disk drives; and
a host computer coupled to the storage server by at least one computer network;
wherein the host computer and the storage server perform input/output (I/O) operations over the at least one network using multiple, concurrent logical
10 connections such that a first I/O operation is executed over a first logical connection while a second I/O operation is executed over a second logical connection.
2. The network storage system as in Claim 1, wherein each logical connection is a socket connection.
- 15 3. The network storage system as in Claim 2, wherein each socket connection is a TCP/IP socket connection.
4. The network storage system as in Claim 1, wherein each logical connection remains persistent over multiple I/O operations performed over that logical connection.
- 20 5. The network storage system as in Claim 1, wherein the host computer is programmed to divide an I/O operation into multiple constituent I/O operations, and to perform the multiple constituent I/O operations in parallel over respective logical connections of said multiple, concurrent logical connections.
- 25 6. The network storage system as in Claim 1, wherein the storage server provides virtualized block-level storage access to the host computer such that the block server appears as local disk drive storage to user-level processes running on the host computer.
- 30 7. The network storage system as in Claim 1, wherein the storage server is configurable to provide multiple storage partitions, each of which may be allocated to a different host computer.

8. The network storage system as in Claim 1, wherein the storage server has a first storage partition which is uniquely assigned to the host computer such that the first storage partition appears to user-level processes running on the host computer as a private, local disk drive.

5 9. The network storage system as in Claim 8, wherein the storage server further has a second storage partition which is uniquely assigned to a second host computer.

10 10. The network storage system as in Claim 8, wherein the storage server further has a second storage partition which is shared by multiple host computers.

10 11. The network storage system as in Claim 1, wherein the host computer and the storage server implement an authentication protocol in which the storage server authenticates the host computer before allowing the host computer to perform input/output operations.

15 12. The network storage system as in Claim 1, wherein the host computer and the storage server implement a discovery protocol in which the storage server notifies the host computer of partitions assigned to the host computer.

13. The network storage system as in Claim 1, wherein at least one of the logical connections is over a general-purpose computer network.

20 14. The network storage system as in Claim 1, wherein at least one of the logical connections is over an Ethernet network.

15. The network storage system as in Claim 1, wherein the first and second logical connections exist over separate computer networks.

16. The network storage system as in Claim 1, wherein each logical connection exists between a respective reader/writer pair.

25 17. The network storage system as in Claim 1, wherein the host computer and the storage server are interconnected by at least one switch.

18. The network storage system as in Claim 1, wherein the host computer and the storage server each include two network interfaces that provide redundant network connections between the host computer and the storage server.

30 19. A system for storing data for host computers, comprising:

a plurality of storage servers connected to a network, each storage server comprising an array of disk drives, an array controller, and a processor;

a plurality of host computers connected to the network and programmed to store data on the storage servers; and

5 at least one switch which interconnects the plurality of storage servers with the plurality of host computers;

wherein each host computer is programmed to open multiple concurrent sockets over the network to the storage servers for performing concurrent input/output operations.

10 20. The system of Claim 19, wherein the sockets are TCP/IP sockets.

21. The system of Claim 19, wherein each storage server of the plurality of storage servers provides virtualized block-level storage access to the host computers such that the block servers appear as local disk drive storage to user-level processes running on the host computers.

15 22. The system of Claim 19, wherein at least a first host computer of the plurality of host computers is programmed to divide an I/O operation into multiple constituent I/O operations, and to perform the multiple constituent I/O operations in parallel over respective logical socket connections between the first host computer and a target storage server.

20 23. The system of Claim 19, wherein a first storage server of the plurality of storage servers is configurable to provide multiple, variable-size partitions, each of which may be allocated to a different host computer of the plurality of host computers.

24. The system of Claim 19, wherein a first storage server of the plurality of storage servers has a first partition which is uniquely assigned to a first host computer of the plurality of host computers such that the first partition appears as a local disk drive to the first host computer.

25 25. The system of Claim 24, wherein the first storage server further has a second partition which is uniquely assigned to a second host computer of the plurality of host computers.

26. The system of Claim 19, wherein the host computers and the storage servers implement an authentication protocol in which a storage server authenticates a host computer before allowing the host computer to perform input/output operations.

27. The system of Claim 19, wherein the host computers and the storage servers implement a discovery protocol in a storage server notifies a host computer of partitions assigned to the host computer.

28. A method of performing input/output operations, comprising:

establishing first and second TCP/IP connections between a host computer and a block-level storage server over one or more computer networks;

performing a first input/output operation over the first TCP/IP connection while performing a second input/output operation over the second TCP/IP connection, each of said input/output operations comprising a transfer of input/output data between the host computer and the storage server; and

maintaining the first and second TCP/IP connections is a persistent state such that each TCP/IP connection may be used to perform additional input/output operations.

29. The method as in Claim 28, wherein the first and second TCP/IP connections are established over separate computer networks.

30. The method as in Claim 28, wherein the first input/output operation is a first I/O request issued from a first application running on the host computer, and the second input/output operation is a second I/O request issued from a second application running on the host computer.

31. The method as in Claim 28, wherein the first and second input/output operations are constituent operations of an I/O request issued by a process running on the host computer, whereby the I/O request is executed in parallel over multiple TCP/IP connections.

32. The method as in Claim 28, further comprising establishing a third TCP/IP connection between the host computer and the storage server, and using the third TCP/IP connection to perform an authentication sequence in which the storage server authenticates the host computer.

33. The method as in Claim 32, further comprising conveying access information over the third TCP/IP connection from the storage server to the host computer, said access information specifying access rights uniquely assigned to the host computer.

5 34. A method of executing an input/output (I/O) request received from a user-level process running on a host computer, comprising:

on the host computer, dividing the I/O request into multiple constituent I/O operations; and

performing the multiple constituent I/O operations in parallel over multiple, respective logical network connections between the host computer and a target storage server such that I/O data is transferred between the host computer and the storage server over each of the logical network connections.

10 35. The method of Claim 34, wherein each logical network connection is a socket connection.

15 36. The method of Claim 34, wherein each logical network connection is a TCP/IP socket connection.

37. The method of Claim 34, wherein at least one of the logical network connections is over a general-purpose computer network.

20 38. A storage server system that provides block-level storage for host computers over a computer network, the system comprising:

a storage server comprising a processor, an array of disk drives coupled to a controller, and a network interface for connecting the storage server to a computer network; and

a software system which provides functionality for subdividing storage space of the array of disk drives into multiple storage partitions, and provides functionality for uniquely allocating the partitions to host computers on the network such that the storage server may be shared by multiple host computers, and such that when a partition is allocated to a host computer, the partition appears to user-level processes of the host computer as a local disk drive.

25 39. The storage server system as in Claim 38, wherein the software system comprises a host driver component which runs on the host computers, and comprises a

server driver component executed by the processor of the storage server, wherein the host and server driver components communicate over the computer network via TCP/IP connections.

5 40. The storage server system of Claim 38, wherein the software system further comprises a configuration and management program which runs on the storage server and provides functionality for partitions to be created and assigned to hosts over the computer network using a web browser.

10 41. The storage server system of Claim 38, wherein the software system further provides functionality for allocating a partition to multiple host computers to permit sharing of partitions.

 42. The storage server system of Claim 38, wherein the partitions have a user-definable size.

15 43. The storage server system of Claim 38, wherein the software system supports the ability for a host computer to concurrently perform multiple input/operations over multiple, respective TCP/IP connections to the storage server.

 44. The storage server system of Claim 38, wherein the software system implements an authentication protocol in which the storage server authenticates host computers using a challenge-response protocol.

20 45. The storage server system of Claim 38, wherein the software system implements a discovery protocol in which the storage server notifies the host computer of storage partitions assigned to the host computer.

 46. A method of providing storage access over a computer network, comprising:

25 establishing a logical connection over a computer network between a host computer and a storage server;

 performing an authentication sequence over the logical connection whereby the storage server authenticates the host computer;

30 when the authentication sequence fails, terminating the logical connection and denying the host computer access to storage resources of the storage server; and

when the authentication sequence is successful, providing access information over the logical connection from the storage server to the host computer to allow the host computer to access the storage resources, at least some of the access information being unique to the host computer.

5 47. The method of Claim 46, wherein the logical connection is a TCP/IP connection.

 48. The method of Claim 46, wherein the authentication sequence comprises the storage server transmitting a randomized challenge to the host computer and thereafter determining whether a host response to the randomized challenge matches an
10 expected response.

 49. The method of Claim 46, wherein the authentication sequence further comprises selecting a software version which is common to the storage server and the host computer, and selecting an authentication method which is dependent upon the selected software version.

15 50. The method of Claim 46, wherein the access information specifies a set of dynamic ports to which the host computer may connect to perform input/output operations.

 51. The method of Claim 46, wherein the access information specifies a maximum number of logical connections that may be established by the host computer
20 to the storage server.

 52. The method of Claim 46, wherein the access information specifies a number of storage partitions allocated to the host computer and sizes of such storage partitions.